Ⓞ

Technical Report 909

AD-A229 106

Performance Utility and Optimal Job Assignment

Roy N. Nord and Leonard A. White

U.S. Army Research Institute

September 1990

DTIC
ELECTE
NOV 1 4 1990
S E D

United States Army Research Institute
for the Behavioral and Social Sciences

FIFTY YEARS OF SERVICE
1940 - 1990

90 11 13 182

U.S. ARMY RESEARCH INSTITUTE

FOR THE BEHAVIORAL AND SOCIAL SCIENCES

A Field Operating Agency Under the Jurisdiction
of the Deputy Chief of Staff for Personnel

EDGAR M. JOHNSON
Technical Director

JON W. BLADES
COL, IN
Commanding

Technical review by

Phil Bobko (Rutgers University)
Curtis Gilroy

Accession For

NTIS GRA&I
DTIC TAB
Unannounced
Justification

By
Distribution/

Availability Codes

| | Avail and/or |
| Dist | Special |

A-1

NOTICES

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS --			
2a. SECURITY CLASSIFICATION AUTHORITY --			3. DISTRIBUTION/AVAILABILITY OF REPORT .Approved for public release; distribution is unlimited.			
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE --						
4. PERFORMING ORGANIZATION REPORT NUMBER(S) ARI Technical Report 909			5. MONITORING ORGANIZATION REPORT NUMBER(S) --			
6a. NAME OF PERFORMING ORGANIZATION U.S. Army Research Institute	6b. OFFICE SYMBOL (If applicable) PERI-RG		7a. NAME OF MONITORING ORGANIZATION --			
6c. ADDRESS (City, State, and ZIP Code) 5001 Eisenhower Avenue Alexandria, VA 22333-5600			7b. ADDRESS (City, State, and ZIP Code) --			
8a. NAME OF FUNDING/SPONSORING ORGANIZATION U.S. Army Research Institute for the Behavioral and Social Sciences	8b. OFFICE SYMBOL (If applicable) PERI-R		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER --			
8c. ADDRESS (City, State, and ZIP Code) 5001 Eisenhower Avenue Alexandria, VA 22333-5600			10. SOURCE OF FUNDING NUMBERS			

10. SOURCE OF FUNDING NUMBERS			
PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO. (211)	WORK UNIT ACCESSION NO.
62785A	791	2101	H1

11. TITLE (Include Security Classification)
Performance Utility and Optimal Job Assignment

12. PERSONAL AUTHOR(S)
Nord, Roy D.; and White, Leonard A.

13a. TYPE OF REPORT Final	13b. TIME COVERED FROM 86/09 TO 90/06	14. DATE OF REPORT (Year, Month, Day) 1990, September	15. PAGE COUNT 37

16. SUPPLEMENTARY NOTATION
--

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	Selection Soldier performance Manpower
05	09		Classification Performance utility
05	08		Optimal assignment Productivity

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

The Army is currently conducting a long-term research project to improve the ability to predict enlisted soldier performance. Also, a new Enlisted Personnel Allocation System (EPAS) is being developed to enhance the effectiveness with which performance information is used to match recruits with Army MOS. To maximize the payoff to the Army of the available manpower pool, the assignment system must consider not only the <u>level</u> at which a recruit will perform in different jobs, but also the <u>utility</u> to the Army of each potential combination of performance level and job. This paper describes a series of simulations undertaken to assess how different models of performance utility will affect the manpower distributions produced by an optimal assignment system. The results show that (a) the use of performance utility in optimal assignment can produce performance gains equal to those achieved when utility is ignored, while simultaneously providing a more balanced distribution of performance across jobs; and (b) specifications in which performance value depends on the mean level of performance seem to produce better results than those that assume constant marginal value.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ DTIC USERS	21 ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL Roy D. Nord	22b TELEPHONE (Include Area Code) (202) 274-5610	22c OFFICE SYMBOL PERI-RG

DD Form 1473, JUN 86   Previous editions are obsolete.   SECURITY CLASSIFICATION OF THIS PAGE

i

Performance Utility and Optimal Job Assignment

Roy N. Nord and Leonard A. White
U.S. Army Research Institute

**Manpower and Personnel Policy Research Group
Curtis L. Gilroy, Chief**

**Selection and Classification Technical Area
Michael G. Rumsey, Chief**

**Manpower and Personnel Research Laboratory
Zita M. Simutis, Director**

U.S. Army Research Institute for the Behavioral and Social Sciences
5001 Eisenhower Avenue, Alexandria, Virginia 22333-5600

Office, Deputy Chief of Staff for Personnel
Department of the Army

September 1990

This document describes research on performance utility conducted as part of a large research effort to improve the selection, classification, and utilization of Army enlisted personnel. The thrust for the project came from the practical, professional, and legal need to validate the Armed Services Vocational Aptitude Battery (ASVAB--the current U.S. military selection and classification test battery) and other selection variables as predictors of training and performance.

The U.S. Army Research Institute for the Behavioral and Social Sciences (ARI) has undertaken a large multiyear research program to meet this requirement and to provide information needed to improve the Army's selection and classification system. One component of the effort referred to as "Project A" is being conducted under contract to the Selection and Classification Technical Area (SCTA) at ARI. The Project A research is aimed at the development, testing, and validation of current and new selection and classification instruments to predict performance in Army training and occupational specialties.

A second key part of this research is the Enlisted Personnel Allocation System (EPAS) being developed by the Manpower Personnel Policy Research Group of ARI with the support of the General Research Corporation, EPAS is a computerized personnel management system that makes extensive use of advanced operations research techniques. The system is capable of incorporating information from Project A on new and revised selection and classification tests to improve the allocation of enlisted personnel. EPAS will improve enlisted personnel performance by achieving a better match between Army job requirements and the capabilities of those applying for service.

Under an optimal selection and classification system the assignment decision would reflect the payoff to the Army of alternative person-job matches. The research described herein compares alternative methods for using performance information to make job assignment decisions using a small-scale prototype of the EPAS.

EDGAR M. JOHNSON
Technical Direc*or

v

PERFORMANCE UTILITY AND OPTIMAL JOB ASSIGNMENT

EXECUTIVE SUMMARY

Requirement:

Each year the Army selects, classifies, and assigns thousands of new recruits to enlisted occupations. To make these personnel decisions in a way that maximizes total system effectiveness, the U.S. Army Research Institute for the Behavioral and Social Sciences has developed an Enlisted Personnel Allocation System (EPAS) to link personnel resources with Army job require- ments. To be fully effective, this system requires not only knowledge of the relationship of personnel aptitudes and characteristics to success on the job, but also the payoff to the Army of alternative assignment decisions.

Research in the Army's Project A has provided information on the utility of different levels of first tour performance in Army jobs. This report describes work to examine the effects of using these job-specific utility functions to make personnel classification and job assignment decisions.

Procedure:

A review of the literature and discussions with Army officers in our utility workshops suggested two models of classification utility. Simulation techniques were used to illustrate the effects of alternative allocation policies on assignment: (a) maximize performance, ignoring performa ce value, and (b) maximize the utility of performance with two different models of utility. In each simulation, a random sample of recruits (1984 accessions) was assigned to nine Army jobs while meeting job manpower demands. The distributions of performance across the nine jobs resulting from each assign- ment strategy were compared to each other and to the expected level of per- formance by soldiers in the sample actually assigned to each job.

Findings:

First, the results clearly indicate that the use of linear optimization techniques to make assignment decisions provides important gains in produc- tivity over those achieved with the current Army classification and job assign-ment system. Second, the policy used to define the payoffs of alternative job assignments has a substantial effect on the expected distribution of perfor-mance across MOS. Assignment to maximize performance produces distributions of performance that are (a) highly variable across jobs, and (b) highly sen-sitive to MOS differences in the validity of the predictor composites and job size.

Relative to current practice, a policy to maximize performance yields the highest productivity gains in jobs for which performance is easiest to measure and predict (i.e., have the highest validities). However, the expected level of performance in jobs with the lowest validities falls below that obtained with the current system.

Two different models of classification utility were defined. Model I assumes that the payoff to an assignment is constant across all assignments to that job and depends only on the expected performance level of the individual selected. Model II, suggested by the economic theory of production, assumes that the marginal utility depends on both the performance level of the individual selected and the average level of performance in the job to which the individual is assigned.

The results show that utility Model II produces different allocations than those resulting from assumption that utility is additive across jobs and individuals. In both models, the effect of maximizing utility rather than performance itself is to weaken the strong link between validity, job size, and assignment. Use of Model II (as compared to Model I) reduces the between-MOS variability in performance to seemingly more acceptable levels, at no cost to performance gains over the current system.

Utilization of Findings:

The results of this work are expected to have two primary uses. First, they will maximize the effectiveness of the newly developed Enlisted Personnel Allocation System (EPAS) as both a policy analysis tool and an operational assignment system. The importance of this role is likely to increase if the improved job-specific performance predictors developed by Project A begin to supplant the current role of Armed Forces Qualification Test scores as the primary measure of recruit quality. Second, the utility analysis results offer an opportunity to substantially improve the precision with which the costs and benefits of manpower policy options can be evaluated.

PERFORMANCE UTILITY AND OPTIMAL JOB ASSIGNMENT

CONTENTS

LIST OF TABLES

CONTENTS (Continued)

PERFORMANCE UTILITY AND OPTIMAL JOB ASSIGNMENT

BACKGROUND

This paper is based on data collected for an Army research project (Project A) aimed at improving the selection and classification system for enlisted personnel. An optimal personnel classification system requires not only information about predicted job performance, but also information about the value to the Army of that performance.

A system that does not explicitly incorporate this information contains an implicit assumption that the value of a given increase in performance is equal across all jobs and at all levels of performance. In economic terms, such an assumption is equivalent assuming that the marginal product of labor is (a) constant and (b) the same in all jobs. Both economic theory and the evidence of current Army policies strongly suggest that such an assumption is unwarranted.

If performance value does vary across jobs and/or levels of performance, then the assumption of constant and equal marginal product can have the perverse effect of producing an allocation that is technically optimal (i.e., one which maximizes expected job performance while meeting all constraints) with a lower value to the Army than random assignment would provide.

To provide the required information, a research effort was undertaken as part of Project A to evaluate the relative payoffs to performance in enlisted occupations. The product of this effort was a set of job-specific "utility functions" that can be used to maximize the gains to the Army of improved classification and job assignment procedures.

Most research on performance utility has addressed the problem of translating performance gains produced by improved selection into a metric that can be used to demonstrate the value of improved selection procedures to skeptical decision-makers. The most common metric is dollar value (e.g., Brogden, 1959; Hunter and Schmidt, 1982), although metrics other than dollar value have also been used (Eaton, Wing, and Mitchell, 1985). In general, the results of this research indicate that substantial productivity gains can result from the use of valid selection procedures.

Estimation of the payoffs to improved classification procedures is a more complex problem (e.g., Brogden, 1959; Schmitz & Nord, 1987). Selection involves choosing the most productive workers for a single job from a pool of applicants. The classification problem assumes that each recruit will be assigned to one of several jobs. Separate equations are used to predict performance in each occupation or job cluster, and the predictions of performance are, in general, highly correlated across jobs. The problem is to assign recruits to jobs in such a way as to maximize productivity while meeting job-specific manpower demands.

In the following section we summarize the methodology employed in the Project A performance utility research, and show the resulting performance utility functions for selected Army MOS. The remainder of the paper describes a series of simulations used to explore the effects of using alternative specifications of the performance utility functions for job assignment.

Performance "Utility" vs. Performance "Value"

Before pursuing the discussion further, a brief digression is in order. The term utility has been widely used in personnel psychology to refer to the value of different levels of job output. This use of the term is appropriate, given the unquantifiable nature of the outputs we are analyzing. It is important, however, to draw the distinction between our use of subjective judgments of performance value and the way similar judgments are used in applications of multi-attribute utility theory, which seek to identify the parameters of individual "utility functions". If we were to interpret the utility judgments we have obtained as reflections of individual preferences, then our use of averages of individual judgments to obtain a single performance value function would require the use of interpersonal comparisons of utility. Such comparisons are invalid under multi-attribute utility theory (e.g., Keeney and Raiffa, 1976). This restriction does not apply to the analysis described here because we do not treat individual judgments of performance value as "utility functions" but rather as imperfect (but randomly distributed) estimates of a single organizational "value function".

MEASURING THE UTILITY OF JOB PERFORMANCE

The Army research on performance utility assessment is being carried out in two stages. The first stage, completed in 1987, focused on the estimation of job-specific performance utility functions for 276 entry-level Army Military Occupational Specialties (MOS). The second stage, which is currently underway, will consist of a series of exercises intended to (a) test the stability of the utility estimates obtained in the first stage; and (b) test the validity of the assumptions required to transform these judgments into aggregate payoff functions that can be used to guide personnel classification and job assignment decisions.

Data Collection

A series of 7 workshops was conducted to obtain judgments of the relative value of performance at five levels in all entry-level Army MOS from 74 field-grade officers. Since the primary focus of this paper is on the assignment effects of utility, we will provide only a very brief summary of the data collection and estimation procedures here. A detailed description can be found in Sadacca, White, Schultz, Campbell, & DiFazio (1989).

The performance level/job combinations were scaled using two methods, one which provided ordinal level utility estimates for 276 MOS, and a second which provided interval-level estimates for twelve of these MOS. Data for seven of the 74 officers in the sample was eliminated because of internal inconsistencies, resulting in a net sample of 67 judges. This sample represented a cross-section of specialties; but the effect of specialty on the utility judgments was generally insignificant.

In the scaling exercises, performance levels were set at the 10th, 30th, 50th, 70th and 90th percentiles, using the current recruit pool as the reference population. Judges were asked to assume that the world was in a state of "heightened tensions". (A brief description of this scenario was provided in the instructions). The "performance" to be evaluated was described as multidimensional, consisting not only of technical proficiency, but also personal discipline and job effort. Judgments were focused on the payoffs to first-tour performance only -- that is, anticipated payoffs to performance at more senior levels were ignored.

The judgments obtained were quite reliable. To eliminate the inflationary effect of the "built-in" agreement that high performance is preferred to low, reliabilities were calculated at each of the five performance percentiles. The resulting n-rater reliabilities ($n=67$) ranged from .82 to .94 with mean .89 for the ordinal judgments, and from .67 to .97 with mean .82 for the interval-level judgments. The inter-method reliabilities at each percentile ranged from .59 to .95, with an average of .77.

Table 1
MOS in Simulation Exercises

MOS	MOS Title	Aptitude Area Composite	Minimum AA Score	1984 Accessions (000's)	Sample N	Minimum Percent I-IIIA
11B	Infantryman	CO	90	17.4	670	61
13B	Cannon Crewmember	FA	85	4.1	158	50
19E	Armor Crewman	CO	90	3.6	137	65
31C	Single Channel Radio Operator	SC	100	.5	20	71
63B	Light Wheel Vehicle Mechanic	MM	90	4.6	176	50
88M	Motor Transport Operator	OF	90	4.8	183	43
71L	Administrative Specialist	CL	95	2.4	94	63
91A	Medical Specialist	ST	95	4.8	183	66
95B	Military Police	ST	100	7.3	281	70

Performance Utility Functions

The 60 performance level/job combinations common to both methods were used to estimate transformation functions from the ordinal to interval scales. The average estimated interval scale values at each performance level were then used to fit a performance utility function for each job. The functions were fitted using stepwise ordinary least squares where the independent variables

were performance level, its square, and its cube. The graphs of these functions for nine MOS are presented in Figure 1. Basic characteristics of these MOS are displayed in Table 1.
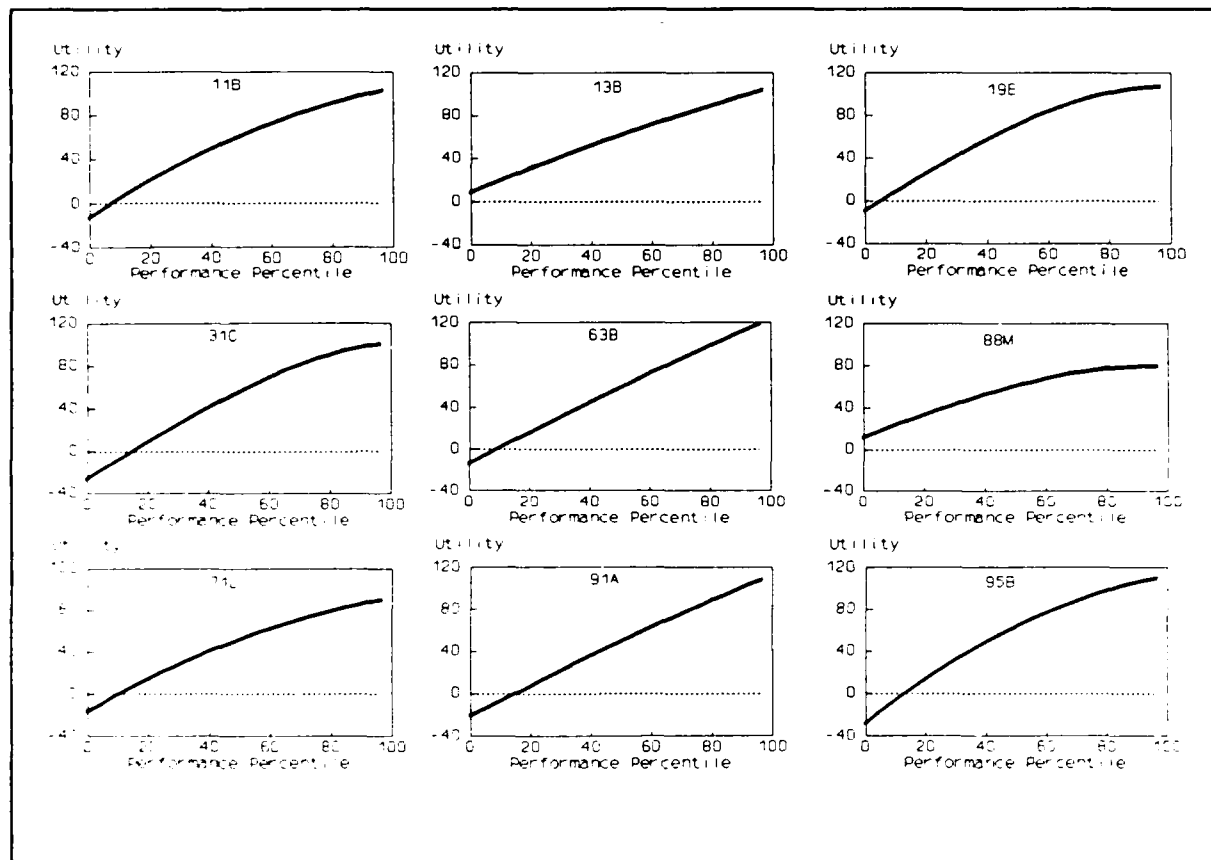
Figure 1. Performance utility functions for nine Army MOS.

These functions illustrate several interesting aspects of our results so far. First, for most MOS the relationship between utility and performance level is a concave function. That is, the functions demonstrate diminishing payoffs to increases in performance as the performance level increases. As we shall see later, this characteristic of the utility function plays a pivotal role in the context of optimal assignment.

A second finding is that there is substantial variety in the shape as well as the intercept (or "scale") of the functions across MOS. One can interpret the scale differences as variations in the "average" value of performance across jobs. In economic terms, this variation can be interpreted as variation across jobs in the marginal product of job output -- that is, differences in the rate at which changes in productivity within a single job contribute to total Army output. Differences in the shape of the functions reflect variations in the way soldier performance at different levels

4

contributes to job output. One would expect, for instance, that functions that are relatively "steep" at low performance levels would be associated with jobs in which the cost of errors is high; and that jobs with relatively steep slopes at high levels of performance would be those in which the payoffs to exceptional performance are high (Bobko & Donnelly, 1986).

On the other hand, as one might expect from previous work in the area of utility generalization (e.g., Bobko, Karren and Kerkar, 1987) there also appear to be identifiable groups of MOS with virtually identical functions. The task of identifying these groups and examining their characteristics is an important subject for further research.

PERFORMANCE UTILITY AND CLASSIFICATION DECISIONS

In this section we address several issues associated with the use of the utility functions developed for each job to make assignment decisions. We begin by examining the consequences of allocating people to jobs so as to maximize predicted performance without taking into account possible variations in performance value. In this analysis, neither job-specific differences in the way manpower contributes to output nor variations in the importance to the Army of the output from different jobs will be reflected in the allocation. We then explore the effects of assignment to maximize utility using two different specifications of the utility functions. Finally, in the concluding section we raise the question of whether the use of performance utility in a classification system will yield better results than would be obtained without it.

Assignment to Maximize Predicted Performance

A simulation was undertaken to illustrate the effects of using a policy that maximizes predicted performance and ignores performance value (or, equivalently, assumes that the marginal value of performance is constant across jobs and levels of performance). The actual distribution of performance for our sample in 1984 is shown in Figure 2.

Each bar in Figure 2 represents the mean performance level of the recruits actually assigned to that job, with performance level measured in standard deviations from the 17-22 year old population mean.

Figure 3 shows a similar representation of the distribution produced by assigning a random sample of recruits (1984 accessions) to nine Army jobs so as to maximize expected performance while meeting job demands (scaled to the sample size in proportion to actual 1984 requirements). All recruits met the minimum entry standards for the jobs to which they were assigned. The optimization used a linear program that maximized the sum across all assignments of predicted performance scores. Predicted performance was calculated using estimated validities of the predictor
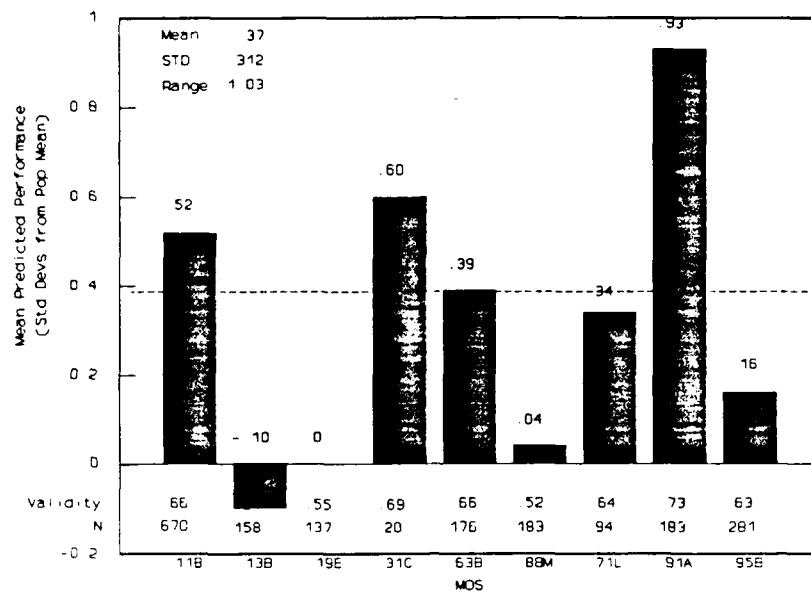
5

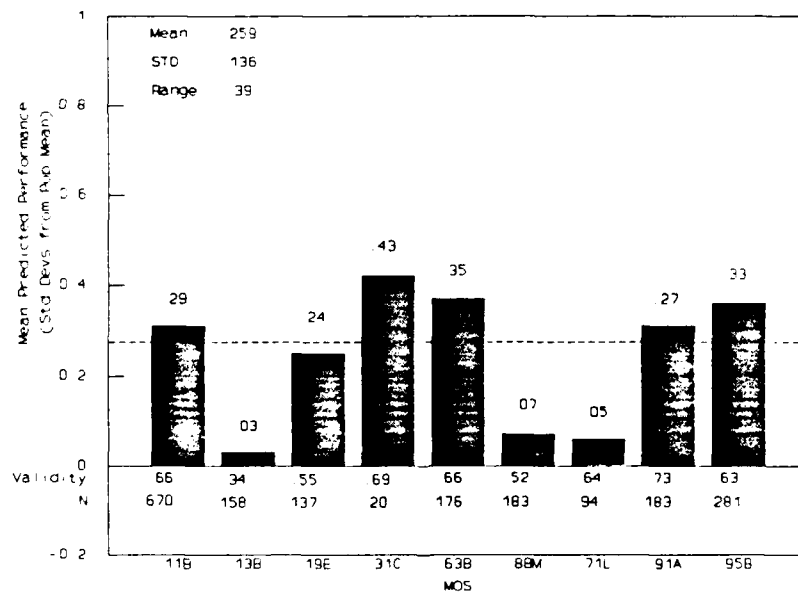Figure 2. Performance distribution when predicted performance is maximized

Figure 3. 1984 actual performance distribution.

composites currently used to predict job performance (Aptitude Area scores) against a combination of multiple choice job knowledge tests and work sample tests of core technical performance.
is maximized. In both figures, the validity and sample size are listed for each job. (Note: The validity and sample sizes are the same for all of the distributions in Figures 2-5. They are included in each figure to make it easier to see the relationships between job quota, validity and the performance distribution across jobs). The mean, standard deviation and range of performance levels across the 9 jobs is also shown on each Figure. Table 2 shows the deviations of the performance levels in each job from the mean across all jobs for the four assignment strategies tested. (Utility Models I and II will be described later.)

Table 2
Summary of Performance Distributions Produced by Four Allocation Strategies

MOS	Actual Assignments	Maximize Performance	Utility Model I	Utility Model II
11B	+.05	+.15	+.08	+.00
13B	-.18	-.47	-.31	-.23
19E	+.05	-.37	-.05	+.08
31C	+.17	+.23	+.23	+.15
63B	+.07	+.01	+.28	+.31
88M	-.19	-.33	-.45	-.31
71L	-.23	-.03	-.24	-.14
91A	-.04	+.56	+.52	+.48
95B	+.15	-.21	-.16	-.10
MEAN	.25	.37	.35	.34
STD	.14	.31	.29	.24

In interpreting these results, it is important to note that two pairs of jobs (MOS 11B & 19E and MOS 91A & 95B) share the same predictor composite (see Table 1) -- that is, they are in the same "job family". This means that predicted performance levels in these pairs will differ only by the difference in validity between the members of the pair. The competition across job families is less direct.

Assignment to maximize predicted performance produced an average MOS performance 0.37 standard deviations above the mean. This represents a substantial gain in productivity over that expected from a policy based on random assignment in which test information is not used. This

gain is quite close to the level of roughly .41[1] predicted by Brogden's (1959) formula for estimating the average productivity of workers assigned by a differential classification battery.

A comparison of Figures 2 and 3 illustrates several consequences of an assignment policy that maximizes performance when the jobs are not of equal size and validity estimates vary by job:

1. Differences in the average level of performance across jobs are substantially increased when predicted performance is maximized. The standard deviation of the mean performance levels shown in Figure 3 is .31 (84% of the mean), as compared with .14 (56% of the mean) for the actual assignments shown in Figure 2.

2. The gains in performance tend to be highest in those jobs for which performance is easiest to measure and predict, and lowest in those jobs with the lowest validities. An ordering of the nine jobs by their validities would yield a very similar list to that produced by ranking average performance levels. The exceptions to this rule are MOS 19E and 88M. (These two jobs use different predictors, and are significantly different in size.)

3. There is a strong interaction effect between validity and job quota. This can be seen by comparing the allocation for MOS 11B to that for 19E. MOS 11B, with a validity of .66 and a (sample) quota of 670, is assigned recruits performing, on average, at .52 standard deviations above the population mean. MOS 19E, which uses the same predictor (CO), has a validity of .55, but a quota only one fifth as large, and receives an allocation performing at the population mean.

Factors Influencing the Effect of Utility on Optimal Assignment

The magnitude of the effects illustrated by the simulation results will depend on a number of environmental and organizational factors. Among these are:

---

[1] Brogden's formula expresses the expected average gain in performance as follows:

Allocation average $= R\sigma(S,K)(1-r)^{1/2}$

where   R is the average validity of the predictors,
r is the average intercorrelation of predicted performance levels across jobs, and
$\sigma(S,K)$ is the expected performance gain that would be observed for rejection ratio (% of applicants rejected) S and number of jobs K if all predictors had a validity of 1.00 and the intercorrelation of predicted performance levels was 0.

The value of .41 was obtained by using R = .602, r = .829 (the average intercorrelation of predicted performance levels in our sample), S = .15, and K = 9. The resulting value of 1.635 for $\sigma(S,K)$ was obtained from Table 1 in Brogden (1959).

This formula relies on a number of assumptions, including multivariate normality of predicted performance in the population, random selection above a specified cutoff point, equal validities across jobs, and natural quotas. Deviations from these assumptions explain the difference between Brogden's prediction and the simulation results.

1. the distribution of the performance predictors in the population;

2. the degree to which performance is differently defined in different jobs (that is, the dimensionality of performance);

3. the variability in validities across jobs and the relationship between validity and job quotas; and

4. the extent to which the allocation is constrained by considerations other than performance.

The effects of 1 and 2 are easiest to explain if we examine them together. If we look at the extremes of the range of these two factors, two effects become clear: If performance is single-dimensioned, or if the predictors of job performance are perfectly correlated in the population, the allocation produced by maximizing expected performance will be exclusively determined by variations across jobs in the predictability of performance. If such variations do not exist, then there will be many equivalent "optimal" allocations.

At the other extreme, if performance is uniquely defined for every job, and the predictors of performance are mutually orthogonal, then the allocation resulting from performance maximization will be unique and identical to the result produced by maximizing any increasing function of performance. In other words, performance utility will be irrelevant to the allocation problem.

With respect to the interaction between validities and job quotas noted in 3, it is obvious that the consequences of variation in predictability will become less pronounced as the variability decreases. Perhaps less obvious is the fact that, if high validities are associated with jobs that have large quotas, the effect of relatively small variations in validity can be exaggerated far out of proportion to the degree of variation. This effect is illustrated in the allocation between MOS 11B and 19E in Figure 3.

Finally, the effect of exogenous constraints (4) is to narrow the range of feasible allocations. The more confining these constraints become, the smaller will be the difference between the "best" and "worst" feasible allocations and thus the smaller the difference induced by considerations of either predicted performance or performance value.

This factor is of particular importance in the case of the Army's allocation problem, which is circumscribed by an extensive set of policy and managerial constraints. These include not only limitations imposed by force structure requirements and the availability of training resources, but also a number of policy constraints whose purpose is to insure an acceptable, if not optimal distribution of performance across jobs. This latter set of constraints includes minimum job entry

standards, an MOS priority system, and a set of job-specific "quality goals" based on educational attainment and scores on the Armed Forces Qualification Test (AFQT). As we shall see, one of the effects of these constraints, when they are used in optimal assignment, is to mitigate the effects of variation in validity and job quotas -- producing an allocation in which average performance is lower, but also less variable across jobs than would occur without them.

If one assumes that these requirements to have evolved in order to enhance Army productivity, then their existence implies two things: (a) that job performance is not equally valuable at all levels in all jobs; and (b) that the payoffs to increases in performance tend to decline in most jobs as the average level of performance increases.

The first conclusion is implied by entry standards, the variation in which is based on the fact that low levels of performance are more tolerable in some jobs than in others. The second is implied by the existence of quality goals, which have two effects: First, differences in goals across MOS imply job-specific differences in the value of high level performance. Second, the role of the goals as constraints in the assignment process has the effect of reducing the payoffs to high-performance assignments in jobs where quality goals are approximately satisfied and increasing these payoffs in jobs that are falling short of the goals.

In the following section, we develop two models of performance utility which deal with these implications of current practice in different ways. After presenting the allocation results produced by these two models, we shall examine the effect on those allocations of adding quality goals as constraints.

Assignment to Maximize Performance Utility

This section considers the use of the information obtained in the utility workshops to make classification and assignment decisions. For these assignment policies the objective is to maximize the utility of performance, rather than performance, per se. In making assignment decisions to maximize utility we face the problem of how to measure the marginal payoff to each assignment; that is, what is the change in the value of job output resulting from a change in the expected level of performance?

A hierarchy of increasingly flexible metrics can be devised. As the flexibility of the metric becomes greater, the functions needed to describe the relationship between performance and output become more complex, and the assumptions about the way job performance contributes to organizational output become less restrictive. There are two key issues here: The first relates to whether or not the value of a given incremental change in performance is constrained to be constant -- i.e., whether or not the utility function is linear. The second involves the way in which performance utility is added up across individuals and across jobs -- that is, the degree of

separability of the function. The simplest approach is to assume (a) that utility is a linear function of the level of performance; and (b) that utility is additive across individuals and jobs.

The most flexible possible approach would allow for the possibility of different non-linear functions for different performance levels and for non-additive aggregation of performance across both individuals and jobs. In the pages to follow, we shall examine the manpower allocations resulting from two specifications that lie between these extremes.

Two Models of Performance Utility

The first model we shall examine (Model I) assumes that performance utility can be added up across individuals and jobs, but makes use of the curvilinear utility functions displayed in Figure 1, thus relaxing the assumption about linear utility functions. Under this model, the payoff of assigning a superior performer to a given job is the same for all assignments to that job, but the gain in utility produced by replacing an average performer with a superior performer will generally be smaller than the gain resulting from a similar exchange between average and inferior performers.

With respect to its assumptions about aggregation of utility across individuals and jobs, Model I is similar to most previously published research on utility (Brogden 1946, 1949; Cronbach & Gleser, 1965; Cascio & Silbey, 1979; Schmidt, Hunter, McKenzie & Muldrow, 1979; Schmidt & Hunter, 1983). Both Model I and these previous studies assume (a) that the total value of performance to an organization can be calculated by measuring the value contributed by individuals performing at different levels in different jobs and simply adding these values across individuals and jobs; and (b) that the relationship between performance level and performance value remains constant as the distribution of performance within the organization changes.

These models differ from ours, however, because they also assume that performance value is a linear function of the level of performance when both value and level are calibrated in standard deviations -- *i.e.*, $\sigma_{value} = \beta \sigma_{performance}$, where σ_{value} is the standard deviation of performance value, $\sigma_{performance}$ is the standard deviation of performance level, and β is the slope of the value function. If both performance level and performance value are assumed to be normally distributed, this formulation implies a shape for the functions shown in Figure 1 exactly the opposite of those we observe -- that is, the higher the level of performance, the steeper the slope of the value function. Figure 4 illustrates this difference for one of the curves shown in Figure 1.

The economic theory of production suggests our second model (Model II), in which the marginal payoff to an assignment depends on both the performance level of the individual and the total "quantity" (or average level) of performance in the job (e.g., Nicholson, 1978). Under this approach, labor is treated as a factor or input to a production process, and the utility or "marginal product" of labor is defined in terms of the change in output resulting from a change in the

11
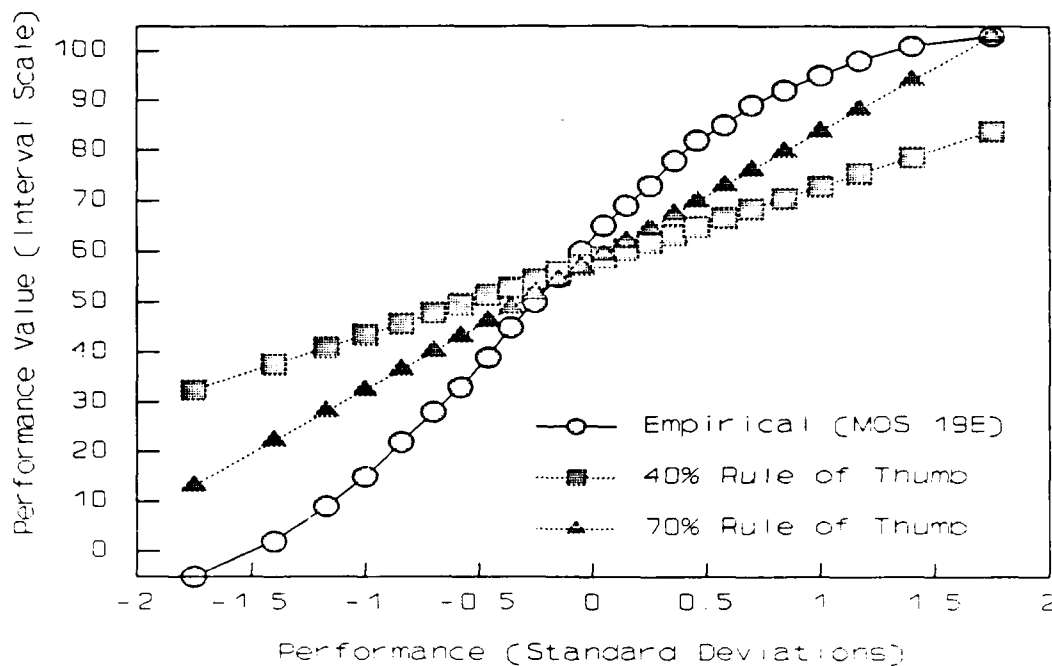
Figure 4. Empirical utility function vs 40% and 70% "rules of thumb" when performance is measured in standard deviations.

quantity of labor, holding all other inputs constant. The economic concept of diminishing returns holds that the marginal product of a given input (job performance in this case) will become smaller as the quantity of that input increases. If this concept applies to job performance then one would predict that the utility functions for job performance would demonstrate the concave shapes evident in the curves shown in Figure 1.

The key difference between Models I and II is that, in Model II, the utility function is applied to the total quantity of performance rather than to individual levels of performance. In order to do this, we must make some assumption about how to aggregate performance across individuals. For this model, we have assumed that this can be done additively -- that is, the aggregate "quantity" of performance is a weighted sum of individual performance, where the weight for a given individual is simply his predicted performance level in percentiles. Thus Model II differs from our first formulation in that it assumes that performance is additive across individuals, but performance utility is not. The same curvilinear utility functions shown in Figure 1 are applied to the aggregate (average) performance level in each job. The difference between the two models in the context of manpower allocation is that, while Model I demonstrates diminishing returns at the level of individual performance, this effect does not apply to aggregate performance -- the value of an additional assignment of a 90th percentile individual to a given job is the same no matter how

12

many 90th percentile individuals have been previously assigned to that job. In Model II, the payoffs depend on the aggregate level of performance.

If it is reasonable to assume that the judgments obtained in the utility workshops are valid when applied to aggregate performance, at least over a limited range, then the generally curvilinear functions displayed in Figure 1 will, under Model II, produce an optimal allocation that is not a "corner solution" -- that is the maximization of performance value will tend to allocate some high-level performers to all jobs. This will occur because of the variations in marginal value implied by the non-constant slopes of the curves. The result is that, for any pair of MOS, X and Y, there will exist some configuration of assignments such that the payoff for an additional high quality assignment to MOS X is greater than that for MOS Y, and some other configuration where the reverse is true.

Note that the assumptions embedded in Model II, while somewhat more flexible than those in the first approach, are nevertheless quite restrictive. Specifically, this model still assumes (a) that the utility of a given level of performance in a given job is independent of the mix of individuals used to achieve that level; and (b) that utility is additive across jobs. Furthermore, in applying the utility judgement data collected so far to aggregate performance, we are taking liberties with the data. The exercises undertaken in the workshops focused on the variations in payoffs to different levels of individual performance in different jobs.

The procedures used in these workshops did not provide direct information on how these payoffs might change as the average level of MOS performance changes. Preliminary results of subsequent workshops conducted to obtain direct estimates of the payoffs to aggregate performance suggest that the aggregate payoff functions will be similar, but not identical, to those used here. However, this evidence is too scanty to provide any firm conclusions regarding the reasonableness of the aggregation assumptions we use here. Thus, the results presented in this section should be interpreted as illustrative.

Throughout the remainder of the paper, we shall occasionally refer to Model I as "strongly separable", and to Model II as "weakly separable". This designation is based on the fact that the utility function for the first model is separable in all of its arguments, while the function for the second model is separable in only some of its arguments. The Appendix provides a mathematical development of the two models and a more explicit definition of "strong" versus "weak" separability.

Utility Maximization Results

Figures 4 and 5, respectively, present the distributions produced when strongly and weakly separable value functions are used to maximize the aggregate utility of performance. These results were obtained using the same sample represented in Figures 2 and 3. The only differences are in

13

Mean Predicted Performance (Std Devs from Pop Mean)

0 8

0 6

0 4

0 2

0

-0 2

35   05   35   58   65   -.06   .06   92   34

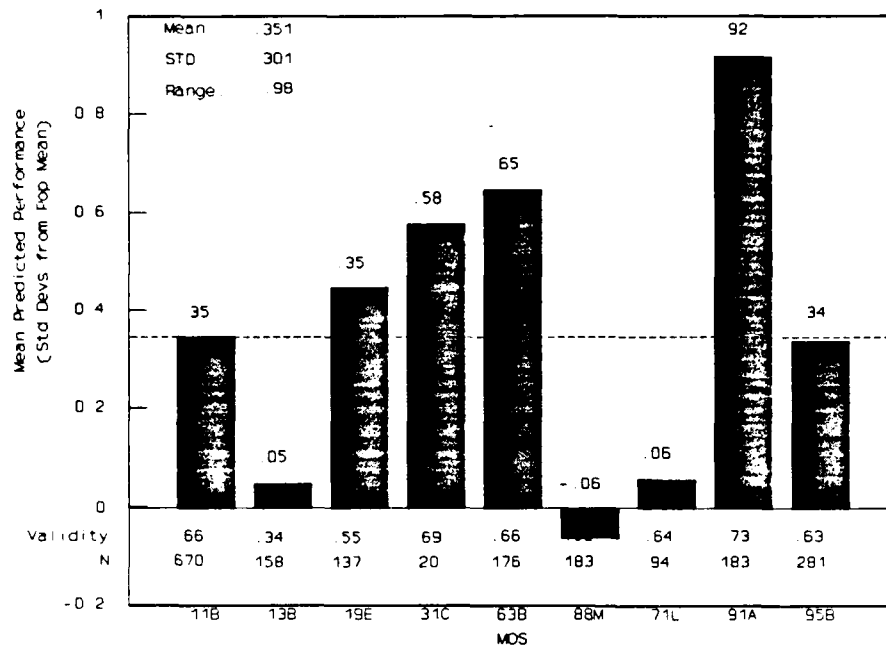| Validity | 66 | .34 | .55 | 69 | .66 | | .64 | 73 | .63 |
| N | 670 | 158 | 137 | 20 | 176 | 183 | 94 | 183 | 291 |
| MOS | 11B | 13B | 19E | 31C | 63B | 88M | 71L | 91A | 95B |

Figure 5. Performance distribution when strongly separable utility is maximized.

the objective functions that were maximized. The results in these Figures are dependent in different ways on how the "tradeoff rates" or marginal rates of substitution vary across performance levels and jobs.

Table 3 provides a list of these rates for three performance levels. Each entry in this table is the ratio of the payoff to an increase in performance in the BASELINE MOS to that for the same increase in the SUBJECT MOS when both MOS are at the baseline performance level. Alternatively, the entries are the ratios of the slopes of the curves in Figure 1 at the baseline percentile levels, where the numerator is the slope of the baseline MOS and the denominator that of the subject MOS.

Note that the single unfavorable ratio for MOS 91A is that involving 63B. Why then does MOS 91A receive a significantly better allocation than 63B? While it is possible that this could happen simply because more recruits had high scores on the predictors of performance in MOS 91A, the more likely reason is the difference in validity between the two MOS. Its higher validity effectively shifts the tradeoff ratios in favor of 91A.

Model I. For this specification, the average level of performance is highest in MOS 91A and 63B. Examination of the rows of Table 3 associated with these two MOS reveals the reason for this. The payoff ratios for 90th percentile performers exceed one in every pair other than 91A vs

14
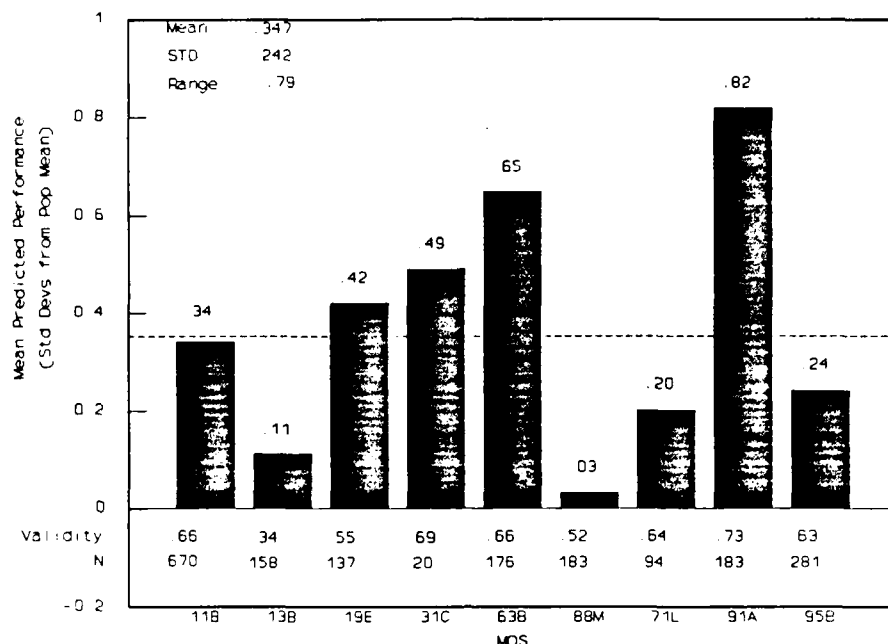
Figure 6. Performance distribution when weakly separable utility is maximized.

63B. Thus, whenever either of these MOS competes with another MOS for an individual performing at a high level in both jobs, the maximum payoff will occur if the individual is assigned to 63B or 91A. In this model, this payoff ratio will apply regardless of the difference in average performance levels between the competing jobs.

Comparison of MOS 31C and 63B, however, show that the effects of validity may be outweighed by the effect of utility. The utility functions for both MOS 31C and 63B have very similar average slopes (Figure 1), but the payoff for high-level performers (see Table 3) is higher for 63B, while the payoffs for low to moderate performance are higher in MOS 31C. This results in a slightly higher average performance level in MOS 63B, as compared to 31C.

The variability of performance across jobs produced by this model is essentially the same as that produced by performance maximization (Table 2), but the variation is less tightly linked to MOS differences in validity. The effect of the interaction between validity and quota evident between MOS 11B and 19E (Figure 3) is also markedly reduced. Conversely, the difference in allocation to MOS 88M and 63B is exaggerated by the use of utility, and the relative positions of MOS 13B and 88M are reversed.

Table 3

Marginal Rates of Substitution Between Pairs of MOS at Three Performance Levels

REFERENCE MOS	LEVEL	11B	13B	SUBJECT 19E	MOS 31C	63B	88M	71L	91A	95B
11B	10	**	1.48	0.97	0.96	1.12	1.59	1.10	1.18	0.81
	50	**	1.20	0.90	0.84	0.85	1.53	1.09	0.87	0.84
	90	**	0.80	2.23	1.20	0.53	7.07	1.06	0.56	0.93
13B	10	0.67	**	0.65	0.64	0.75	1.07	0.74	0.79	0.54
	50	0.84	**	0.75	0.70	0.71	1.28	0.91	0.72	0.70
	90	1.24	**	2.77	1.49	0.66	8.80	1.33	0.70	1.16
19E	10	1.04	1.54	**	0.99	1.16	1.65	1.13	1.22	0.84
	50	1.11	1.33	**	0.94	0.95	1.71	1.21	0.96	0.94
	90	0.45	0.36	**	0.54	0.24	3.17	0.48	0.25	0.42
31C	10	1.05	1.55	1.01	**	1.17	1.67	1.15	1.23	0.84
	50	1.19	1.42	1.07	**	1.02	1.82	1.29	1.03	**
	90	0.84	0.67	1.86	**	0.44	5.91	0.89	0.47	0.78
63B	10	0.89	1.33	0.86	0.85	**	1.42	0.98	1.05	0.72
	50	1.17	1.40	1.05	0.98	**	1.80	1.27	1.01	0.98
	90	1.88	1.51	4.19	2.25	**	13.29	2.00	1.05	1.75
88M	10	0.63	0.93	0.61	0.60	0.70	**	0.69	0.74	0.51
	50	0.65	0.78	0.59	0.55	0.56	**	0.71	0.57	0.55
	90	0.14	0.11	0.32	0.17	0.08	**	0.15	0.08	0.13
71L	10	0.91	1.35	0.88	0.87	1.02	1.45	**	1.07	0.74
	50	0.92	1.10	0.83	0.77	0.78	1.41	**	0.80	0.77
	90	0.94	0.75	2.09	1.12	0.50	6.64	**	0.53	0.88
91A	10	0.85	1.26	0.82	0.81	0.95	1.35	0.93	**	0.69
	50	1.16	1.38	1.04	0.97	0.99	1.77	1.26	**	0.97
	90	1.78	1.43	3.98	2.13	0.95	12.62	1.90	**	1.67
95B	10	1.24	1.84	1.20	1.18	1.39	1.97	1.36	1.46	**
	50	1.19	1.43	1.07	1.00	1.02	1.82	1.29	1.03	**
	90	1.07	0.86	2.39	1.28	0.57	7.58	1.14	0.60	**

Note in Table 3 that MOS 88M faces unfavorable tradeoff ratios against every other job at every performance level -- i.e., it will always be "last in line". The exaggeration of the difference between MOS 63B and 88M occurs because this effect is exacerbated by the lower validity in 88M. MOS 13B provides an interesting contrast. In spite of its extremely flat utility function (Figure 1), Table 3 reveals that this MOS can compete on a favorable basis for 90th percentile performers with MOS 11B, 19E, 31C, 88M, and 71L, indeed, it is only the low validity of its predictor that prevents this MOS from receiving a dramatically increased allocation.

Model II. Figure 5 displays the results under Model II (weakly separable utility). Inter-job variability in average performance is less than that produced by either performance maximization or

16

Table 4

Payoff Intersection Points at Three Performance Levels Under Weakly Separable Utility Model

REFERENCE MOS	LEVEL	11B	13B	SUBJECT 19E	31C	MOS 63B	88M	71L	91A	95B
11B	30	**	+	43	48	34	+	19	12	48
	50	**	+	57	63	-	+	42	93	63
	70	**	66	68	75	-	35	64	-	78
13B	30	59	**	62	69	-	09	52	-	70
	50	65	**	66	73	-	27	59	-	74
	70	71	**	69	76	-	37	65	-	79
19E	30	17	+	**	34	+	+	05	+	38
	50	39	+	**	56	72	+	30	64	55
	70	74	78	**	77	-	40	68	-	81
31C	30	13	+	26	**	+	+	01	+	30
	50	32	+	45	**	43	+	22	34	50
	70	61	35	63	**	-	16	54	-	71
63B	30	29	+	42	47	**	+	18	-	48
	50	34	+	46	51	**	+	24	43	51
	70	39	+	50	55	**	+	29	62	55
88M	30	67	56	67	74	-	**	61	-	70
	50	82	-	74	82	-	**	78	-	87
	70	-	-	84	93	-	**	-	-	-
71L	30	39	-	50	56	72	+	**	64	55
	50	57	24	61	68	-	+	**	-	69
	70	75	84	71	78	-	42	**	-	82
91A	30	32	+	44	50	41	+	21	**	50
	50	36	+	47	53	56	+	26	**	52
	70	41	+	51	57	80	+	32	**	57
95B	30	05	+	10	15	+	+	+	+	**
	50	32	+	45	50	43	+	22	30	**
	70	59	30	62	69	-	10	52	-	**

Model I, and the resulting distribution bears a greater resemblance to the actual distribution than do the other optimizations. However, the gains over the current system in overall average performance produced by both Model I and performance maximization are not lost -- this allocation achieves an average gain of .09 standard deviations over the actual assignments.

Several other differences between Model II and the actual distributions are noteworthy:

(a) The greatest relative increases in performance levels occur in MOS 63B and 91A, which both show large proportional increases over their actual levels. All of these jobs tend to have steep linear utility functions.

(b) The converse of this tendency can be seen by examining the jobs which show the greatest relative <u>declines</u> in average performance levels. These tend to be those with the most pronounced curvature in their utility functions -- MOS 31C, 88M, and 95B. Pairwise comparisons of MOS 31C with 63B and 91A with 95B reveal reversals of the relative allocations to the pairs in each case.

Comparing these results to those shown in Figure 3 (performance maximization), we can see that the extreme effect of the interaction between validity and quota evident between MOS 11B and 19E is considerably weakened. In Model II, MOS 19E is allocated an average performance level roughly 25% higher than that assigned to 11B.

Comparing Figures 4 and 5, we can see the differences between the allocations produced by Models I and II. The most noticeable of these is the reversal of the relative positions of MOS 11B and 19E. The tendency of the weakly separable function to "even out" the distribution is also evident in the fact that, in all but one case, MOS that are below the mean in Figure 4 receive increased levels of performance in Figure 5, while the reverse is true for those jobs allocated above-average levels under Model I.

To help analyze the information displayed in Figure 5, we have constructed Table 4 to provide information about marginal rates of substitution between pairs of MOS when the performance levels in the members of the pair are at different levels. Table 4 provides this information for three "baseline levels" of performance.

In this table, we set the average performance level in a "reference MOS" at one of three baseline levels (30th, 50th or 70th percentile), and ask the following question: At what level of performance in each other ("subject") MOS will the payoff to an increment of performance be the same as that in the comparison MOS. In other words, if we fix one element of a pair of performance levels at the baseline, what is the second element of the pair that will yield a marginal rate of substitution at unity?

The numbers (30, 50, or 70) along the left-hand side of the table indicate "baseline" levels of average performance for the indicated <u>reference</u> MOS (listed in the far left column). The numbers in the body of the table are the levels of average performance in the <u>subject</u> MOS (listed across the top of the Table) at which the payoff to a small increase in performance in the reference MOS is equal to the payoff for the same increase above the baseline level for the subject MOS. For all levels below this point, the payoff will be greater for the <u>reference</u> MOS, and for all levels above this point, the payoff will be greater for the <u>subject</u> MOS. Entries coded with a "+" indicate that the payoff to for the subject MOS at the baseline is <u>larger</u> than that for the reference MOS no matter what the performance level in the reference MOS. A "-" entry indicates the reverse -- the payoff to the subject MOS at the baseline is always <u>less</u> than that to the reference MOS. The "**" entries

18

denote the intersections between each MOS and itself, which are by definition equal to the baseline levels.

For example, if we examine the first three rows (where MOS 11B is the reference MOS) of the columns associated with MOS 13B and 91A, we can conclude the following:

1. When 11B has an average performance level below 50 the payoffs for performance gains in 11B will be higher than those in 13B, no matter what the average level in 13B. However, by the time that average performance in 11B has reached the 70th percentile, this is no longer true. At this point, the payoffs will favor 13B as long as it's average performance level is below the 66th percentile.

2. Looking at the column for MOS 91A, the table shows that, when 11B is at an average level of 30 and 91A is at or above the 12th percentile, the payoffs will favor 11B. However, as the 11B level rises above 30, the differential declines rapidly. By the time 11B reaches an average of 50, 91A offers higher payoffs as long as it's average level is below the 93rd percentile. By the time the level in 11B reaches 70, the payoffs will favor 91A, even if it has already been assigned a pool with average performance in the 99th percentile.

We can use this information to see why the reversal of MOS 11B and 19E occurs. The table shows that the slopes of the utility functions for MOS 11B and 19E are roughly equal when both have average performance levels at the 68th percentile. Below this level the payoffs are higher for 19E and above it they are higher for 11B. In the strongly separable case this means that, given a choice, the optimization will "prefer" to assign all individuals below the 60th percentile to 19E and those above that level to 11B. The effect of this "preference" (combined with the higher validity in 11B) is to produce a higher average performance level in 11B.

In the weakly separable case, the optimization's "decisions" are governed by marginal rates of substitution that change as the average levels of performance in the two jobs changes. Table 4 shows that when MOS 19E has an average performance level at the 50th percentile, its payoffs will equal or exceed those to 11B, as long as 11B has an average level below the 39th percentile. As the average level in 19E increases toward 70, this differential becomes smaller. Beyond approximately the 68th percentile, the increasing curvature in the utility function for MOS 19E (Figure 1) causes the payoff ratio to shift gradually in favor of 11B. Thus, in order for MOS 11B to receive an average performance level higher than 19E, both MOS would have to receive allocations above the 68th percentile. Given the available population, the requirements of 11B, and the competition from other MOS, such an allocation is not feasible. (Translated into percentile terms, the levels shown in Figure 5 are 52 and 66 for 11B and 19E, respectively.) The result, relative to that in Model I, is a substantial reallocation of performance away from MOS 11B in favor of 19E.

Optimal Assignment and AFQT-based Quality Goals

We noted earlier that the current Army job assignment system includes a set of AFQT-based "quality goals". These goals are specified as minimum percentages of AFQT category I-IIIA accessions in each job. (There are also goals in the form of caps on IIIB and IV accessions, but we shall focus here on the effect of the high-quality goals only.) These goals have inter-related effects: First, because of the generally strong correlation between AFQT and job performance, the goals tend to produce a more balanced distribution of first-tour soldier performance across jobs than would otherwise occur. Second, the goals ensure that every MOS has a pool of soldiers with strong general aptitude from which to "grow" its non-commissioned officers (NCOs) in the second tour and beyond. The first of these two effects is similar to that embedded in the weakly separable utility model. The second effect is not addressed in the utility work undertaken so far. To incorporate this factor into the utility model will require both the development of effective predictors of NCO performance, and extension of the utility assessment to include NCO performance.

In order to assess the effect of including quality goals, we added the quality goals as constraints to each of three optimizations described above. The results of this exercise are displayed in Figure 6. (The 1984 I-IIIA goals for each MOS are displayed in Table 1.) The dotted lines in Figure 6 show the mean 1984 performance level, and the solid lines indicate the mean level for each of the optimal assignments.

In interpreting these results, it is important to keep in mind two important factors affecting the impact of quality goals on the performance distribution. First, in 1984, the goals were "tight" -- that is, virtually 100% of AFQT I-IIIA accessions were required to meet the MOS targets. In 1986, the goals were only 95% of high-quality accessions, and in 1987, this percentage was lower still. The quality goals enter the optimization as constraints, meaning that the optimization pays no attention to performance (or performance value) until all quality goals are met. Thus, the smaller the difference between quality requirements and total high-quality accessions, the less "room" is available to maximize performance value. A second consideration is the degree to the predictors of job performance used in the optimization are correlated with AFQT. The higher this correlation, the smaller will be the effect of differences in the way performance utility is modelled. The average correlation in our sample between AFQT and predictor scores is roughly .8. We would expect this correlation to diminish somewhat in the future as a result of ongoing Project A efforts to improve criterion measures and develop better job-specific performance predictors.

The results shown in Figure 6 have three significant implications:

1. In spite of the tightness of the quality goals and the relatively high correlation between AFQT and our predictor set, all of the optimal assignments show substantial gains in
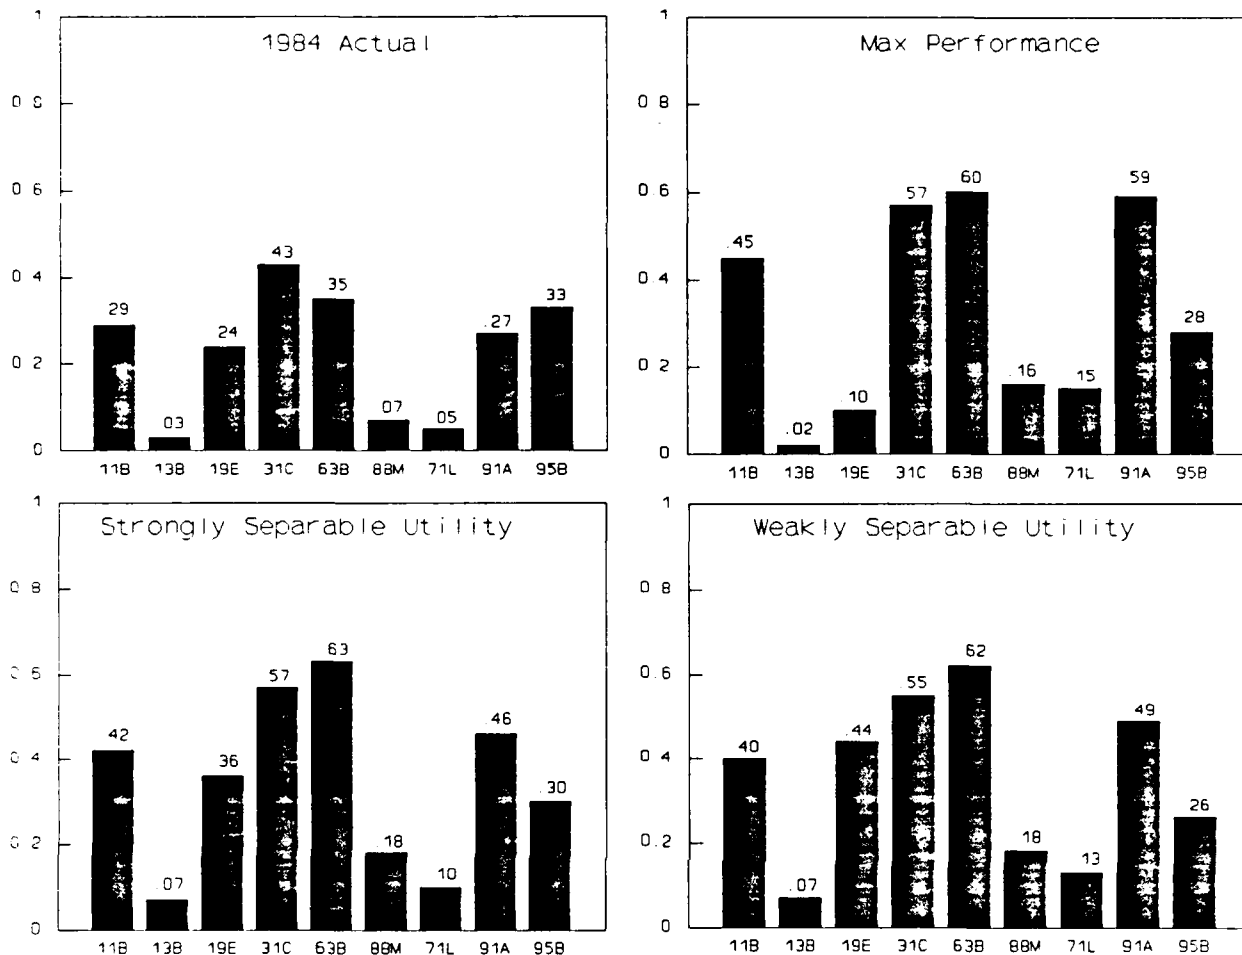
20

Figure 7. Performance distributions under alternative utility models when AFQT quality goals are enforced.

average performance over the current system. While some of these gains are due to the fact that the optimizations do not reflect real-world constraints (such as time-specific training requirements, sequential assignment, applicant choice, etc.), previous simulations of the Enlisted Personnel Allocation System (EPAS) have shown that a substantial portion of these gains can be retained even when these constraints are imposed (Schmitz and Nord, 1987).

2. The introduction of quality goals erases the differences in the distributions produced by utility Models I and II. Note, however, that either a change in the quality of the

accession pool or a reduction in the correlation between AFQT and job performance would again produce differences similar to those shown in Figures 4 and 5.

3. The effect of the quota/validity interaction between MOS 11B and 19E continues to be highly evident in the performance maximization model. The two utility models produce distributions that are much more likely to be seen as "sensible" by Army decision-makers.

SUMMARY AND CONCLUSION

Several aspects of the results presented are noteworthy. First, the job-specific utility functions revealed substantial variation in the payoffs to different levels of performance in Army jobs. Second, the results pictured in Figures 2-5 provide ample evidence that the approach used to define the expected payoffs to classification decisions can make considerable difference in the resulting manpower allocations.

For the assignment policies we examined, the use of linear optimization techniques provided significant performance gains over those achieved under the current system. However, assignment to maximize performance produces distributions that are highly variable across jobs, and highly sensitive to the interaction between job size and validity. When a policy to maximize performance was used, performance allocated to several jobs with relatively lower validities fell below the level provided by the current assignment system. Utility maximization under a strongly separable value function results in a weakening of the link between validity, job size, and assignment; but does not significantly reduce inter-job variation in average performance levels.

The results demonstrate that a utility model which assumes a weakly separable value function produces different allocations than those resulting from direct application of the same functions to individual performance. The weakly separable approach reduces inter-job variability in performance while providing significant increases in average performance over the current system. This model thus provides an opportunity to make maximum use of both current and enhanced predictors of job performance in a way that is responsive to Army objectives other than individual job performance. Further research is needed to determine whether or not the allocation will be similarly sensitive to variations in the assumptions we have used to arrive at a measure of aggregate performance.

The differences between the two types of utility models are diminished when AFQT-based quality goals are introduced, but even when quality goals are "tight" and AFQT-performance correlations are high, it appears likely that some consideration of performance utility will be essential in order to make effective use of optimal assignment in an operational context. As to whether it will be possible to replace the use of quality goals with a model that relies only on a

weakly separable utility function, the answer will depend on (a) the addition of second-tour performance to the utility functions; (b) the degree to which Army policy makers perceive job performance as distinct from ability as measured by AFQT; and (c) the extent to which the procedures used to arrive at estimates of performance utility are accepted as legitimate by those policy-makers.

One problem that awaits clearer resolution is that of defining the terms upon which alternative approaches are to be compared. It does seem clear, however, that the objective of a selection and classification system is not simply to maximize job performance per se, but to maximize the productivity of the human resources available to the organization. To do this, it is essential that information on the value of performance be incorporated into the selection and classification system. Furthermore, the evidence provided by current practice strongly indicates that, in the judgement of Army decision-makers, (a) personnel allocation decisions should not be driven by individual performance alone; and (b) the tradeoffs among different performance level/job combinations are not constant.

REFERENCES

Bobko, P., & Donnelly, L. (1986). Toward an understanding of overall worth in utility analyses: A policy capturing approach. Manuscript submitted for publication.

Bobko, P., Karren, R., & Kerkar, S. P. (1987). Systematic research needs for understanding supervisory-based estimates of SDy in utility analyses. Organizational Behavior and Human Decision Processes, 69, 69-95.

Brogden, H. (1959). Efficiency of classification as a function of number of jobs, percent rejected, and the validity and intercorrelation of job performance estimates. Educational and Psychological Measurement, 9, 181-190.

Eaton, N. K., Wing, H., & Mitchell, K. J. (1985). Alternate methods of estimating the dollar value of performance. Personnel Psychology, 38, 27-40.

Hunter, J. & Schmidt, F. (1982). Fitting people to jobs: the impact of personnel selection on national productivity. In Human Performance and Productivity: Human Capability Assessment. Hillsdale New Jersey: Erlbaum.

Keeney, R. L. & Raiffa, H. (1976). Decisions with multiple objectives: Preferences and value tradeoffs. New York: Wiley.

Nicholson, Walter. (1978). Micro Economic Theory: Basic Principles and Extensions (2nd ed.). The Dryden Press: Hinsdale, IL.

Sadacca, R., & Campbell, J. P. (1985). Assessing the utility of a personnel classification system. Paper presented at the meeting of the Southwestern Psychological Association, Atlanta, GA.

Sadacca, R., White, L. A., Schultz, S. R., Campbell, J. C., & DiFazio, A. S. (1989). Assessing the utility of MOS performance levels in Army enlisted occupations. (ARI Technical Report 839). Alexandria, VA: US Army Research Institute. (AD A211 608)

Schmidt, F., Hunter, J. E., Outerbridge, A. N., & Trattner, M. H. (1986). The economic impact of job selection methods on size, productivity, and payroll costs of the Federal work force: An empirically-based demonstration. Personnel Psychology, 39, 1-29.

Schmidt, F. L., Hunter, J. E., & Pearlman, K. (1982). Assessing the economic impact of personnel programs on workforce productivity. Personnel Psychology, 15, 333-348.

Schmitz, E., & Nord, R. (1987). Evaluating improvements to MOS assignment policy. (MPPRG Working Paper). Alexandria, VA: US Army Research Institute.

APPENDIX

An Analytic Representation of the Allocation Problem

The allocation problem can be described as follows: Let N be the total number of positions to be filled, M be the number of jobs, and K the number of levels of performance. We can then represent any assignment of N individuals to M jobs by an M x K matrix $\mathbf{Q}$, where q_{ij} is the number of individuals at performance level j assigned to job i. If we define a k x 1 vector $\mathbf{p}$ such that p_i is the quantity of performance obtained from an individual performing at level i (the elements of $\mathbf{p}$ might be performance percentiles, for instance), then we can define a scalar Z, the total quantity of performance represented by the allocation $\mathbf{Q}$ as

$$Z = \mathbf{p'Q} \tag{1}$$

That is, the total quantity of performance represented by the allocation $\mathbf{Q}$ is simply the sum of the number of individuals assigned to each job, weighted by performance level. This is the definition of aggregate performance that we will use. However, before continuing, it should be noted that such a definition implicitly assumes that the total quantity of performance obtained is independent of how performance is distributed within and across jobs. In other words, we are ignoring issues relating to unit or group performance.

Given this definition of aggregate performance, we must define a way of applying a performance utility function to the quantity Z, that is we must define a function v(Z) using the payoff functions obtained for each job. We shall consider two alternative specifications:

(a) <u>Model I</u> assumes that v(Z) is a "<u>strongly separable</u>" function of $\mathbf{p}$ and $\mathbf{Q}$ that can be written in the form

$$v(Z) = \sum_{i=1}^{M} \sum_{j=1}^{K} q_{ij} u_j(p_i) \tag{2}$$

where $u_j(p_i)$ is the value of performance at level i in job j.

If we assume strong separability, the gain in utility from any single assignment depends solely on the performance level of the assignment being considered. In this model, the payoff is independent of the mean level of performance assigned to the job. That this must be true can be shown by simply differentiating (2) with respect to q_{ij}, as follows:

$$\frac{\partial v(Z)}{\partial q_{ij}} = u_j(p_i), \text{ for } 0 \leq q_{ij} \leq N, \text{ i}\in K, \text{ j}\in M. \tag{3}$$

(b) <u>Model II</u> assumes that v(Z) is "<u>weakly separable</u>"[1] -- that is

$$v(Z) = \sum_{i=1}^{M} u_j(q_i p), \text{ where } q_i \text{ is the jth row of } \mathbf{Q}. \tag{4}$$

By relaxing the separability assumption, we allow the marginal value of an additional assignment to a given job to vary with the total <u>quantity</u> of performance in that job as well as with the performance <u>level</u> of the particular assignment being considered:

$$\frac{\partial v(Z)}{\partial q_{i,j}} = h_j(\mathbf{p'q'}) \tag{5}$$

The consequences choosing Model I rather than Model II for optimal assignment can be seen by comparing the maximization problems associated with the two specifications.

Let d_j represent the demand (quota) for job j, and s_i be the supply of applicants (recruits) predicted to perform at level i. (For now, we assume that performance is unidimensional -- that is, each applicant will perform at the same level in all jobs.) Then the performance value functions defined by (2) and (4), produce the following optimal assignment problems:

<u>Maximize</u>

Model I: $\quad \sum_{i=1}^{M} \sum_{j=1}^{K} q_{i,j} u_j(p_i)$ $\tag{6}$

or

Model II: $\quad \sum_{i=1}^{M} u_j(\mathbf{q_i p})$ $\tag{7}$

These objective functions are maximized subject to demand and supply constraints:

Demands: $\quad \sum_{i=1}^{M} q_{i,j} = d_j, \text{ for all } j \in M$ $\tag{8}$

Supplies: $\quad \sum_{j=1}^{K} q_{i,j} = s_i, \text{ for all } i \in K$ $\tag{9}$

---

[1] More precisely, this alternative assumes that v(Z) is separable in <u>jobs</u>, but not in <u>performance levels</u>. Mathematically, this implies

$$\partial \left[\frac{\partial u(.)/\partial q_{\cdot j}}{\partial u(.)/\partial q_{\cdot k}} \right] \Big/ \partial q_{\cdot m} = 0 \quad \underline{\text{but}} \quad \partial \left[\frac{\partial u(.)/\partial q_{rj}}{\partial u(.)/\partial q_{\cdot j}} \right] \Big/ \partial q_{tj} \neq 0,$$

where r,s and t index performance levels, and j,k, and m index jobs.

The equation systems defined by (a) and (b) can be transformed into single equations using the method of Lagrange as follows:

Model I:
$$L = \sum_{i=1}^{M} \sum_{j=1}^{K} q_{ij} u_j(p_i) + \sum_{j=1}^{K} \gamma_j \sum_{i=1}^{M} (q_{ij} - d_j) + \sum_{i=1}^{M} \pi_i \sum_{j=1}^{K} (q_{ij} - s_i)$$ (10)

or

Model II:
$$L = \sum_{i=1}^{M} \sum_{j=1}^{K} u_j(p'q_j') + \sum_{j=1}^{K} \gamma_j \sum_{i=1}^{M} (q_{ij} - d_j) + \sum_{i=1}^{M} \pi_i \sum_{j=1}^{K} (q_{ij} - s_i)$$ (11)

where γ_j and π_i are sets of Lagrangian multipliers associated with the demand and supply constraints.

The conditions for a maximum of (10) will be easier to describe if we order the values of $u_j(p_i)$ so that the following is true:

> If $j' > j$ then $u_{j'}(p_i) \le u_j(p_i)$
> and if $i' > i$ then $u_j(p_{i'}) \le u_j(p_i)$

Then the matrix of assignments $\mathbf{Q}^*$ that maximizes (a) will contain elements q_{ij} that meet the following condition:

$$q_{ij} = \text{MAX} \ \{s_i - \sum_{k=1}^{i-1} q_{kj}, \ d_j - \sum_{l=1}^{k-1} q_{im}\}.$$ (12)

In other words, the maximum will be achieved by following the simple rule of "top-down" assignment: Order the set of possible person-job matches from those with the highest value to those with the lowest; then assign individuals at the highest available level of performance to the position with the highest value at that level of performance until either the demand is met or the supply is exhausted.

The necessary (first order) conditions for a $\mathbf{Q}^*$ that maximizes (b), the weakly separable case, can be stated as follows:

$\mathbf{Q}^* = \{q_j^*\}$ such that

(i) $\quad \dfrac{\partial L}{\partial q_{ij}} \bigg|_{q_{ij} \in q_j^*} = \dfrac{\partial u_j(q_j^* p)}{\partial q_{ij}} \bigg|_{q_{ij} \in q_j^*} \cdot \gamma_i - \pi_j = 0$, for all i,j (13)

(ii) $\quad \dfrac{\partial L}{\partial \gamma_i} \bigg|_{q_{ij} \in q_j^*} = s_i - \sum_{j=1}^{K} q_{ij} = 0$, for all i (14)

(iii) $\quad \dfrac{\partial L}{\partial \pi_i} \bigg|_{q_{ij} \in q_j^*} = d_j - \sum_{i=1}^{M} q_{ij} = 0$, for all j (15)

The solution of this system implies that, if the functions u_j are continuous, twice differentiable, and convex, there will exist a unique optimal solution that is characterized by the following:

$$\frac{\partial u_j(.)/\partial q_{ij}}{\partial u_k(.)/\partial q_{ik}} = \frac{\partial u_j(.)/\partial q_{mj}}{\partial u_k(.)/\partial q_{mk}} \qquad \text{for all } i \neq m, j \neq k, \qquad (16)$$

and

$$\frac{\partial u_j(.)/\partial q_{ij}}{\partial u_j(.)/\partial q_{mj}} = \frac{\partial u_k(.)/\partial q_{ik}}{\partial u_k(.)/\partial q_{mk}} \qquad \text{for all } i \neq m, j \neq k, \qquad (17)$$

That is, at optimality, the marginal rates of substitution across jobs for the same performance level will be the same for all pairs of jobs and performance levels, as will the marginal rates of substitution among different performance levels within jobs.